arXiv:1403.7058v1 [cs.DS] 27 Mar 2014

The Sketching Complexity of Graph Cuts

Alexandr Andoni* Microsoft Research
Robert Krauthgamer† Weizmann Institute
David P. Woodruff‡ IBM Almaden

March 28, 2014

Abstract

We study the problem of sketching an input graph, so that, given the sketch, one can estimate the value (capacity) of any cut in the graph up to $1+\varepsilon$ approximation. Our results include both upper and lower bound on the sketch size, expressed in terms of the vertex-set size n and the accuracy ε.

We design a randomized scheme which, given $\varepsilon \in (0,1)$ and an n-vertex graph $G=(V,E)$ with edge capacities, produces a sketch of size $\tilde{O}(n/\varepsilon)$ bits, from which the capacity of any cut $(S, V \setminus S)$ can be reported, with high probability, within approximation factor $(1+\varepsilon)$. The previous upper bound is $\tilde{O}(n/\varepsilon^2)$ bits, which follows by storing a cut sparsifier graph as constructed by Benczúr and Karger [BK96] and followup work [SS11, BSS12, FHHP11, KP12].

In contrast, we show that if a sketch succeeds in estimating the capacity of all cuts $(S, \bar{S})$ in the graph (simultaneously), it must be of size $\Omega(n/\varepsilon^2)$ bits.

1 Introduction

In 1996 Benczúr and Karger [BK96] introduced *cut sparsifiers*, a remarkable and very influential notion: given a graph $G=(V,E,w)$ with $n=|V|$ vertices, $m=|E|$ edges and edge weights $w: E \to \mathbb{R}_+$, together with a desired error parameter $\varepsilon \geq 0$, the cut sparsifier is a sparse graph H on the same n vertices (in fact a subgraph of G but with different edge weights), such that every cut in G is $(1+\varepsilon)$-approximated by the corresponding cut in H. Specifically, they show there always exists such a sparsifier H with $O(\frac{n}{\varepsilon^2}\log n)$ edges only, potentially much less than in the original graph G, and that it can be constructed in time $O(m\log^2 n)$. This construction has proven to have tremendous impact on cut problems in graphs, see e.g. [BK96, BK02, KL02, She09, Mad10].

Naturally, the notion has been studied further, leading to improved bounds [GRV09, FHHP11, KP12], as well as to the generalization of the notion, termed *spectral sparsifiers* [ST04, ST11]. A spectral sparsifier is a graph H such that the quadratic form associated with the Laplacian of H provides a $1+\varepsilon$ approximation to the quadratic form associated with the Laplacian of G. The original bound of $O(\frac{n}{\varepsilon^2}\log^{O(1)} n)$ was later improved to $O(\frac{n}{\varepsilon^2}\log n)$ [SS11], and finally to $O(\frac{n}{\varepsilon^2})$ by Batson, Spielman, and Srivastava [BSS12]. The latter result improves also the original construction

*Email: andoni@microsoft.com

†Work supported in part by a US-Israel BSF grant #2010418, an Israel Science Foundation grant #897/13, and by the Citi Foundation. Email: robert.krauthgamer@weizmann.ac.il

‡Email: dpwoodru@us.ibm.com

of Benczúr and Karger. These spectral sparsifiers have been instrumental in obtaining the first near-linear time algorithm for solving SDD linear systems [ST04], and as such have led to many beautiful ideas down the road, see also [KMP10, KMP11].

It is natural to ask whether one can improve the dependence on ε in the size of H. The Alon–Boppana theorem and [Alo97] suggest the answer is no: if H is constrained to be a regular graph (unweighted, up to scaling), then $\Omega(n/\varepsilon^2)$ edges are required to approximate all the cut values of the complete graph (see also [Nil91, BSS12]). Note that for the complete graph, a matching upper bound can be achieved by regular expander graphs [LPS88, MSS13].

We show that one can obtain a cut sparsification with *linear* dependence on $1/\varepsilon$, if some constraints are relaxed. Specifically, we construct a randomized *sketch* (succinct data structure) of size $O(\frac{n}{\varepsilon}(\log n)^{O(1)})$, that given any query $S \subset V$, outputs the weight of the respective cut $w(S,\bar{S}) := \sum_{e\in E:|e\cap S|=1} w(e)$, approximated within $1+\epsilon$ with high probability. Notice that our result relaxes the usual notion of cut sparsification in two ways: (1) it is a sketch rather than a graph; and (2) its guarantee is "with high probability for each cut" rather than "for all cuts" (in the compressed sensing literature, these guarantees are often termed "for each" and "for all", respectively).

Are these relaxations necessary? We show that at least the second one, i.e., the "for each" guarantee, is indeed necessary. Specifically, we prove that any sketch that, for every given graph, approximates *all* of its cuts within approximation factor $1+\varepsilon$ must have size $\Omega(n/\varepsilon^2)$ bits. Thus, our lower bound strengthens the Alon–Boppana and [Alo97] lower bounds in the following sense. Our lower bound implies that if every graph has a not necessarily regular, and possibly weighted, cut sparsifier graph with s edges, then $s = \Omega(n/(\varepsilon^2 \log n))$.

Moreover, our lower bound shows for any (possibly randomized) sketching algorithm, if the requirement is only that for each input graph G on n vertices, with constant probability the algorithm produces a sketch approximating all cuts of G, then the sketch size is $\Omega(n/\varepsilon^2)$ bits. This shows there is a large family of "hard" graphs on which any sketching algorithm with smaller sketch size will fail. Our $\Omega(n/\varepsilon^2)$ bit lower bound also applies to sketches which approximate the quadratic form of G up to $1+\varepsilon$, as such sketches also approximate all cuts of G. This generalizes an $\Omega(n/\varepsilon^2)$ edge lower bound of [BSS12] for spectral sparsifiers of the complete graph, as ours holds for all sketching algorithms, not only those which approximate G by another graph H.

There are several streaming algorithms which produce cut or spectral sparsifiers which achieve $\tilde{O}(n/\varepsilon^2)$ bits of space [AG09, KL13, AGM12, GKK10]. Our lower bound shows these algorithms are optimal, up to logarithmic factors.

Finally, we argue that, even if we lose on some guarantees by going from a graph sparsifier to a sketch, the sketch is nonetheless useful algorithmically. For example, our results immediately yield a 2-pass streaming algorithm $1+\varepsilon$-approximating the min-cut of the graph, using space $\tilde{O}(n/\varepsilon)$; see Appendix C for a short proof. All previous streaming results for this problem [AG09, KL13, AGM12, GKK10] had a quadratic dependence on $1/\varepsilon$ in space.

We remark that our sketch size is optimal up to polylogarithmic size. In particular, it is not possible to obtain a "for each" sketch with size less than $O(n/\varepsilon)$ (see Appendix D for a proof).

1.1 Techniques

Upper Bound. We now discuss the main ideas behind our sketch construction. Let us first give some intuition why the previous algorithms cannot yield a $\tilde{O}(n/\varepsilon)$ bound, and show how our

algorithm circumvents these roadblocks on a couple of illustrative examples. For concreteness, it is convenient to think of $\varepsilon = 1/\sqrt{n}$.

All existing cut and spectral sparsifiers construct the sparsifier graph by taking a subgraph of the original graph G, with the "right" re-weightening of the edges [BK96, SS11, BSS12, FHHP11, KP12]. In fact, except for [BSS12], they all proceed by sampling edges independently, each with its own probability (that depends on the graph).

Consider for illustration the complete graph. In this case, these sampling schemes employ a uniform probability $p \approx \frac{1/\varepsilon^2}{n}$ of sampling every edge. It is not hard to see that one cannot sample edges with probability less than p, as otherwise anti-concentration results suggest that even the degree of a vertex (i.e., the cut of a "singleton") is not preserved within $1+\varepsilon$ approximation. Perhaps a more interesting example is a random graph $\mathcal{G}_{n,1/2}$, where edges are sampled independently with (roughly) uniform probability, that again cannot be less than p, because of singleton cuts. However, if we aim for a sketch for the complete or $\mathcal{G}_{n,1/2}$ graph, we can just store the degree of each vertex using only $O(n)$ space, and this will allow us to report the value of every singleton cut (which is the most interesting case, as the standard deviation for these cut values is of the order of $1 + \varepsilon = 1 + 1/\sqrt{n}$). These observations suggest that *sketching* a graph may go beyond considering a subgraph (or a different graph) to represent the original graph G.

Our general algorithm proceeds in several steps. The core of our algorithm is a procedure for handling cuts of value $\approx 1/\varepsilon^2$ in a graph with unweighted edges, which proceeds as follows. First, repeatedly partition the graph along every *sparse* cut, namely, any cut whose sparsity is below $1/\varepsilon$. This results with a partition of the vertices into some number of parts. We store the cross-edges (edge connecting different parts) explicitly. We show the number of such edges is only $\tilde{O}(n/\varepsilon)$, and hence they fit into the alloted space. Obviously, the contribution of these edges to any desired cut $w(S, \bar{S})$ is easy to compute.

The final component of the sketching algorithm thus estimates the contribution (to a cut $w(S, \bar{S})$ for a yet unknown $S \subset V$) from edges that are inside any single part P of the partition. To accomplish this, we sample $\approx 1/\varepsilon$ edges out of each vertex, and also store the exact degrees of all vertices. Then, to estimate the contribution of edges inside a part P to $w(S, \bar{S})$, we take the sum of (exact) degrees of all vertices in $S \cap P$, *minus* an estimate for (twice) the number of edges inside $S \cap P$ (estimated from the edge sample). This "difference-based" estimate has a smaller variance than a direct estimate for the number edges in $(S \cap P, \bar{S} \cap P)$ (which would be the "standard estimate", in some sense employed by previous work). The smaller variance is achieved thanks to the facts that (1) the assumed cut is of size (at most) $1/\varepsilon^2$; and (2) there are no sparse cuts in P.

Overall, we achieve a sketch size of $\tilde{O}(n/\varepsilon)$. We can construct the sketch in polynomial time by employing an $O(\sqrt{\log n})$-approximation algorithm for sparse cut [ARV09, She09] or faster algorithms with $(\log^{O(1)} n)$-approximation [Mad10].

The main upper bound theorem is Theorem 2.1.

Lower Bound. We prove our lower bound using communication complexity. The natural thing to do would be to give Alice a graph G, and Bob a cut S. Alice produces a sketch of G and sends it to Bob, who must approximate the capacity of S. The communication cost of this problem lower bounds the sketch size. However, as we just saw, Alice has an upper bound with only $\tilde{O}(n/\varepsilon)$ bits of communication. We thus need for Bob to solve a much harder problem which uses the fact that Alice's sketch preserves all cuts.

We let G be a disjoint union of $\varepsilon^2 n/2$ graphs G_i, where each G_i is a bipartite graph with $\frac{1}{\varepsilon^2}$

vertices in each part. Each vertex in the left part is independently connected to a random subset of half the vertices in the right part. Bob's problem is now, given a vertex v in the left part of one of the G_i, as well as a subset T of half of the vertices in the right part of that G_i, decide if $|N(v) \cap T| > \frac{1}{4\varepsilon^2} + \frac{c}{\varepsilon}$, or if $|N(v) \cap T| < \frac{1}{4\varepsilon^2} - \frac{c}{\varepsilon}$, for a small constant $c > 0$. Most vertices v will satisfy one of these conditions, by anti-concentration of the binomial distribution. Note that this problem is not a cut query problem, and so a priori it is not clear how Bob can use Alice's sketch to solve it.

To solve the problem, Bob will do an exhaustive enumeration on cut queries, and here is where we use that Alice's sketch preserves all cuts. Namely, for each subset S of half of the vertices in the left part of G_i, Bob queries the cut $S \cup T$. As Bob ranges over all (exponentially many) such cuts, what will happen is that for most vertices u in the left part for which $|N(u) \cap T| > \frac{1}{4\varepsilon^2} + \frac{c}{\varepsilon}$, the capacity of $S \cup T$ is a "little bit" larger if u is excluded from S. This little bit is not enough to be detected, since $|N(u) \cap T| = \Theta\left(\frac{1}{\varepsilon^2}\right)$ while the capacity of $S \cup T$ is $\Theta\left(\frac{1}{\varepsilon^4}\right)$. However, as Bob range over all such S, he will eventually get lucky in that S contains all vertices u for which $|N(u) \cap T| > \frac{1}{4\varepsilon^2} + \frac{c}{\varepsilon}$, and now since there are about $\frac{1}{2\varepsilon^2}$ such vertices, the little $\frac{c}{\varepsilon}$ bit gets "amplified" by a factor of $\frac{1}{2\varepsilon^2}$, which is just enough to be detected by a $(1+\varepsilon)$-approximation to the capacity of $S \cup T$. If Bob finds the S which maximizes the (approximate) cut value $S \cup T$, he can check if his v is in S, and this gives him a correct answer with large constant probability.

Our lower bound theorem is Theorem 3.1.

2 Sketching Algorithm for Graph Cuts

Let $G = (V, E, w)$ be a graph with $n = |V|$ vertices and edge weights $w : E \to \mathbb{R}_+$. The goal is to construct a sketch for G that preserves each cut with constant probability. For this goal we prove the following theorem:

Theorem 2.1. *Given a weighted graph $G = (V, E, w)$ on n vertices, where the non-zero weights are in the range $[1, W]$, and $1/n \le \varepsilon \le 1/30$, there exists a cut sketch of size $\tilde{O}(n/\varepsilon \cdot \log\log W)$ bits. Specifically, for every query $S \subset V$, the sketch produces a $1 + O(\varepsilon)$ approximation to $w(S, \bar{S})$, with probability at least $7/9$.*

In this section, we focus on the case when the non-zero edge weights are in the polynomial range, i.e., $W = n^5$, as this is the crux of the construction. We show how to extend the construction to the general–weights case in Appendix B.

We will say that $\hat{a} > 0$ is a ρ-approximation to $a > 0$ (for $\rho \ge 1$) if their ratio is $\hat{a}/a \in [1/\rho, \rho]$. For an edge weights function w, let $w(A, B)$ denote the total weight of edges connecting two disjoint subsets $A, B \subset V$.

2.1 Construction

Below, we show how we sketch the graph, as well as how we estimate the cut size given the sketch.

Sketching Algorithm. The sketch has two components. The first component is a standard 1.4-cut sparsifier (recall that a $(1+\varepsilon)$-cut sparsifier is a sparse graph H on same vertex-set, which approximates every cut in G within factor $1+\varepsilon$). We can use the construction of Benczúr and Karger [BK96], or subsequent constructions [SS11, BSS12, FHHP11, KP12] (some of which produce

a spectral sparsifier, which is only stronger); any of these methods will produce a graph H with $\tilde{O}(n)$ edges.

The second component is the main ingredient of the sketch, and is described next. Let $\tilde{C} = \{1.4^i \mid 0 \le i \le \log_{1.4} n^5\}$ be the set of size $O(\log n)$ such that each cut value in G is 1.4-approximated by some value $c \in \tilde{C}$. For each value $c \in \tilde{C}$ we construct a structure D_c as follows. First by scaling all the edge weights, let us assume $c = 1$. Now discard every edge e whose (scaled) weight is $w_e > 5$.

In the next step, called importance sampling, we sample each (remaining) edge $e \in E$ independently with probability $p_e := \min\{w_e/\varepsilon^2, 1\}$, and assign the sampled edges new edge weights $\tilde{w}_e := w_e/p_e$. Notice that $\tilde{w}_e \in [\varepsilon^2, 5]$. (It may be convenient to consider the non-sampled edges as having weight $\tilde{w}_e := 0$.) Now let $\tilde{E}$ be the set of sampled edges, and partition it into $l = O(\log \frac{1}{\varepsilon})$ classes according to the (new) edge weights, namely, $\tilde{E} = L_1 \cup \cdots \cup L_l$ where $L_i = \{e \in E : \ \tilde{w}_e \in (5 \cdot 2^{-i}, 5 \cdot 2^{-i+1}]\}$.

For each class L_i, recursively partition the graph (V, L_i) as follows: break the current cluster $P \subseteq V$ whenever it contains a subset $P' \subset P$ of size $|P'| \le |P|/2$ such that $d_i(P', P \setminus P')/|P'| \le 1/\varepsilon$, where $d_i(A, B)$ is the number (not weight) of class L_i edges connecting between two disjoint subsets $A, B \subset V$. Once the recursive partitioning process is finished, denote the resulting partition of V by $\mathcal{P}_i$, and store in the sketch all the edges $E_{\mathcal{P}_i}$ connecting different parts of $\mathcal{P}_i$. In addition, store for every vertex $v \in V$ its "weighted degree", i.e., the total weight of its incident class L_i edges $\tilde{w}_i(v) := \sum_{e \in L_i : v \in e} \tilde{w}_e$. From the remaining edges (those inside a part of $\mathcal{P}_i$), we store the following sample: for every non-isolated vertex store $s = 1/\varepsilon$ incident edges, each chosen uniformly at random with replacement. For every stored edge $e \in L_i$ (from $E_{\mathcal{P}_i}$ or not), the sketch keeps the edge weight $\tilde{w}_e$.

Estimation algorithm. Given a query subset $S \subset V$, first use the graph sparsifier H to compute $\tilde{c}$, a 1.4-approximation to the desired cut value $w(S, \bar{S})$, and use the structure D_c, where $c \in \tilde{C}$ is a 1.4-approximation to $\tilde{c}/(1.4)^2$, namely, $c \in [\tilde{c}/(1.4)^3, \tilde{c}/1.4]$. Now estimate the contribution to the cut from edges in each class L_i as follows.

The contribution from a single class L_i, $i \in [l]$ is composed of two terms. The first term is the total $\tilde{w}_e$ weight of edges between S and $\bar{S}$ that are in the set $E_{\mathcal{P}_i}$ (recall the sketch stores all edges between different parts of $\mathcal{P}_i$).

The second term estimates the $\tilde{w}_e$ weight of edges inside each part $P \in \mathcal{P}_i$. Specifically, check whether $|S \cap P| \le |P|/2$ or $|\bar{S} \cap P| < |P|/2$. In the first case, the estimate I_P for P is given by the sum over all vertices in $S \cap P$, of their "weighted degree" (which is stored in the sketch) minus the (appropriately scaled) weight $\tilde{w}$ of sampled edges that are inside $S \cap P$. Formally, we define

$$I_P := \sum_{x \in S \cap P} \Big[\tilde{w}_i(x) - \tfrac{d_{ix}(P)}{s} \sum_{e \in L_{ix}} 1_{\{e \in E(S, \bar{S})\}} \tilde{w}_e \Big], \tag{1}$$

where $d_{ix}(T)$ is the number of class L_i edges from vertex x to a set $T \subset V$, and L_{ix} is the multiset of s class L_i edges incident to x (the sample chosen by the sketching algorithm). In the second case ($|\bar{S} \cap P| < |P|/2$), the estimate is similar except that we now use the "weighted degrees" in $\bar{S} \cap P$ and total weight of edges inside $\bar{S} \cap P$.

The overall second term is the sum of these estimates over all parts $P \in \mathcal{P}_i$. The final estimate is just the sum of these two terms over all classes L_i.

2.2 Size and Correctness Guarantees

First we show the bound on space usage. The sparsifier H has $\tilde{O}(n)$ edges. By construction, we have $O(\log n)$ possible cut values $\tilde{C}$ and for each one, we have $l = O(\log \frac{1}{\varepsilon}) \leq O(\log n)$ edge weight classes. For each weight class L_i, the sketch stores (1) at most $O(\frac{1}{\varepsilon} n \log n)$ edges in $E_{\mathcal{P}_i}$, because each step in the recursive partitioning process contributes $d_i(P', P \setminus P')/|P'| \leq 1/\varepsilon$ edges per vertex in P', and each vertex appears in the smaller subset P' at most $\log n$ times; and (2) at most n/ε sampled edges, because for every non-isolated vertex we sample $s = 1/\varepsilon$ incident edges. All in all, we have $\tilde{O}(n/\varepsilon)$ edges, each requiring $O(\log n)$ bits.

We proceed to analyze the accuracy of the estimation procedure. Fix a query $S \subset V$. After consulting the sparsifier H, we can approximate $w(S, \bar{S})$ within factor $(1.4)^2 < 2$ and thus use a data structure D_c where $c \leq w(S, \bar{S}) \leq 4c$. Thus, by rescaling to $c = 1$, we need to $1 + \varepsilon$ estimate the cut value $w(S, \bar{S})$ which is between 1 and 4.

First, note that the discarded edges do not affect the solution since they could not have been part of $w(S, \bar{S})$ as they are too heavy. Second, we bound the effect of the importance sampling step on the value of the cut, namely, that with high probability $\tilde{w}(S, \bar{S})$ is a $1 + \varepsilon$ approximation to $w(S, \bar{S})$. Indeed, its expectation $\mathbb{E}[\tilde{w}(S, \bar{S})] = w(S, \bar{S})$ since every edge e that was not discarded satisfies $\mathbb{E}[\tilde{w}_e] = w_e$ (and more generally, it is a Horvitz-Thompson estimator), and its variance is

$$\mathrm{Var}[\tilde{w}(S, \bar{S})] = \sum_{e \in E(S,\bar{S})} \mathrm{Var}[\tilde{w}_e] = \sum_{e \in E(S,\bar{S})} w_e^2/p_e - w_e^2 \leq \sum_{e \in E(S,\bar{S})} \varepsilon^2 w_e = \varepsilon^2 w(S, \bar{S}),$$

where the inequality is verified separately for $p_e = 1$ and for $p_e = w_e/\varepsilon^2$. Thus, by Markov's inequality, with probability at least 8/9, we have $|\tilde{w}(S, \bar{S}) - w(S, \bar{S})| \leq 3\varepsilon\sqrt{w(S, \bar{S})} \leq 3\varepsilon \cdot w(S, \bar{S})$.

Next, we show that the rest of the procedure is likely to estimate $\tilde{w}(S, \bar{S})$ well. Let c_i be the contribution to $\tilde{w}(S, \bar{S})$ by edges in class L_i; hence $\sum_{i=1}^{l} c_i = \tilde{w}(S, \bar{S})$ and each $c_i \leq \tilde{w}(S, \bar{S}) \leq (1+3\varepsilon)\, w(S, \bar{S}) \leq 5$. Let us also denote $\lambda_i := 5 \cdot 2^{-i}$, then all edges $e \in L_i$ have weight $\tilde{w}_e \in [\lambda_i, 2\lambda_i]$. Let $\hat{c}_i$ be the contribution of edges of class L_i to the estimate computed by the algorithm; it is equal to $\sum_{e \in E_{\mathcal{P}_i} \cap E(S,\bar{S})} \tilde{w}_e$, which is computed exactly because the sketch stores all edges between different parts of $\mathcal{P}_i$, plus the sum over all parts $P \in \mathcal{P}_i$ of their estimate I_P. Recall that, in (1), each I_P is the sum of "weighted degrees" in either $S \cap P$ or $\bar{S} \cap P$ (whichever has smaller cardinality), minus the appropriately scaled weight of edges inside that subset.

It is easy to verify that $\hat{c}_i$ is an unbiased estimator for c_i, namely, $\mathbb{E}[\hat{c}_i] = c_i$. We now analyze its variance, which comes only from the estimators I_P of edges inside each $P \in \mathcal{P}_i$. We can assume that for each $P \in \mathcal{P}$ we have $|S \cap P| \leq |P|/2$ (otherwise, exchange $S \cap P$ and $\bar{S} \cap P$). Let $c_i(P)$ be the weight of the cut $\tilde{w}(S, \bar{S})$ inside P, restricted to edges from class L_i. Then $\sum_{P \in \mathcal{P}_i} c_i(P) \leq c_i$. In each part P, the variance of I_P comes only from the sampled edges (since the "weighted degrees"

are known exactly):

$$\begin{aligned}
\mathrm{Var}[I_P] &= \sum_{x\in S\cap P}\Big(\tfrac{d_{ix}(P)}{s}\Big)^2 \mathrm{Var}\Big[\sum_{e\in L_{ix}} 1_{\{e\in E(S,\bar{S})\}}\tilde{w}_e\Big] \\
&\le \sum_{x\in S\cap P}\Big(\tfrac{d_{ix}(P)}{s}\Big)^2 \cdot s\cdot \tfrac{d_{ix}(S\cap P)}{d_{ix}(P)}\cdot (2\lambda_i)^2 \\
&= 4\varepsilon\lambda_i^2 \sum_{x\in S\cap P} d_{ix}(P)\cdot d_{ix}(S\cap P) \\
&\le 4\varepsilon\lambda_i^2\cdot |S\cap P| \sum_{x\in S\cap P} d_{ix}(P) \\
&\le 4\varepsilon\lambda_i^2\cdot |S\cap P|\cdot \Big[d_i(S\cap P,\bar{S}\cap P) + 2|S\cap P|^2\Big].
\end{aligned}$$

By the stopping condition of the recursive partitioning $|S\cap P| \le \varepsilon\cdot d_i(S\cap P,\bar{S}\cap P)$, and thus

$$\mathrm{Var}[I_P] \le 4\varepsilon^2\lambda_i^2 (d_i(S\cap P,\bar{S}\cap P))^2\Big[1+2\varepsilon^2 d_i(S\cap P,\bar{S}\cap P)\Big].$$

Observe that $c_i(P)$ consists of $d_i(S\cap P,\bar{S}\cap P)$ edges whose $\tilde{w}$ weight is in the range $[\lambda_i, 2\lambda_i]$, hence $\lambda_i d_i(S\cap P,\bar{S}\cap P) \le c_i(P)$. Using $c_i(P)\le c_i\le 5$ and $\lambda_i\ge\varepsilon^2$, we can further derive $\varepsilon^2 d_i(S\cap P,\bar{S}\cap P)\le 5$, which together give

$$\mathrm{Var}[I_P] \le 44\varepsilon^2 (c_i(P))^2.$$

Therefore, the total variance over all parts $P\in\mathcal{P}_i$ is

$$\mathrm{Var}[\hat{c}_i] = \sum_{P\in\mathcal{P}_i}\mathrm{Var}[I_P] \le 44\varepsilon^2\sum_{P\in\mathcal{P}_i}(c_i(P))^2 \le 44\varepsilon^2 c_i^2.$$

It follows that the algorithm's final estimate $\hat{c} := \sum_{i=1}^{l}\hat{c}_i$ is unbiased, namely,

$$\mathbb{E}[\hat{c}] = \sum_i \mathbb{E}[\hat{c}_i] = \sum_i c_i = \tilde{w}(S,\bar{S}),$$

and its total variance (over all levels i) is

$$\mathrm{Var}[\hat{c}] = \sum_i \mathrm{Var}[\hat{c}_i] \le 44\varepsilon^2\sum_i c_i^2 \le 44\varepsilon^2\tilde{w}(S,\bar{S}).$$

Thus, by Markov's inequality, with probability at least $8/9$, we have that $|\hat{c}-\tilde{w}(S,\bar{S}))| \le 3\varepsilon\sqrt{44\,\tilde{w}(S,\bar{S})} \le 21\varepsilon\,\tilde{w}(S,\bar{S})$, where the last inequality uses that $\tilde{w}(S,\bar{S})\ge(1-3\varepsilon)\,w(S,\bar{S})\ge 9/10$. Altogether, we obtain that with probability at least $7/9$, the algorithm's estimate $\hat{c}$ is a $1+O(\varepsilon)$ approximation for $\tilde{w}(S,\bar{S})$, and also the latter quantity is a $1+O(\varepsilon)$ approximation for $w(S,\bar{S})$, Or, more directly,

$$|\hat{c}-w(S,\bar{S})| \le 3\varepsilon\,w(S,\bar{S}) + 21\varepsilon(1+3\varepsilon)\,w(S,\bar{S}) \le 27\varepsilon\,w(S,\bar{S}),$$

which completes the proof of the theorem.

3 Lower Bound

We prove the following lower bound on the sketch size of graphs.

Theorem 3.1. *Fix an integer n and $\varepsilon \in [\frac{1}{n}, 1]$, and let $\mathbf{sk} = \mathbf{sk}_{n,\varepsilon}$ and $\mathbf{est} = \mathbf{est}_{n,\varepsilon}$ be "sketching" and "estimation" algorithms for (unweighted) graphs on vertex set $[n]$, which may be randomized. Suppose that for every such graph $G = ([n], E)$, with probability at least $3/4$ the resulting sketch $\mathbf{sk}(G)$ satisfies*[1]

$$\forall S \subset [n], \quad \mathbf{est}\left(S, \mathbf{sk}(G)\right) \in \ (1 \pm \varepsilon) \cdot |E(S, \bar{S})|.$$

Then the worst-case space requirement of $\mathbf{sk}$ is $\Omega(n/\varepsilon^2)$ bits.

Our proof of this theorem uses the following communication lower bound for a version of the Gap Hamming Distance problem, which we prove in Section 3.1. Throughout, we fix $c := 10^{-3}$ (or a smaller positive constant), and assume $\varepsilon \leq c/10$.

Theorem 3.2. *Consider the following distributional communication problem: Alice has as input n strings $s_1, \ldots, s_{n/2} \in \{0,1\}^{1/\varepsilon^2}$ of Hamming weight $\frac{1}{2\varepsilon^2}$, and Bob has an index $i \in [n/2]$ together with one string $t \in \{0,1\}^{1/\varepsilon^2}$ of Hamming weight $\frac{1}{2\varepsilon^2}$, drawn as follows:*[2]

- *i is chosen uniform at random;*
- *s_i and t are chosen uniformly at random but conditioned on their Hamming distance $\Delta(s_i, t)$ being, with equal probability, either $\geq \frac{1}{2\varepsilon^2} + \frac{c}{\varepsilon}$ or $\leq \frac{1}{2\varepsilon^2} - \frac{c}{\varepsilon}$;*
- *the remaining strings $s_{i'}$ for $i' \neq i$ are chosen uniformly at random.*

Consider a (possibly randomized) one-way protocol, in which Alice sends to Bob an m-bit message, and then Bob determines, with success probability at least $2/3$, whether $\Delta(s_i, t)$ is $\geq \frac{1}{2\varepsilon^2} + \frac{c}{\varepsilon}$ or $\leq \frac{1}{2\varepsilon^2} - \frac{c}{\varepsilon}$. Then Alice's message size is $m \geq \Omega(n/\varepsilon^2)$ bits.

We can interpret the lower bound of Theorem 3.2 as follows: Consider a (possibly randomized) algorithm that produces an m-bit sketch of Alice's input $(s_1, \ldots, s_{n/2}) \in \{0,1\}^{n/\varepsilon^2}$, and suppose that the promise about $\Delta(s_i, t)$ can be decided correctly (with probability at least $3/4$) given (only) the sketch and Bob's input $(i, t) \in [n/2] \times \{0,1\}^{1/\varepsilon^2}$. Then $m \geq \Omega(n/\varepsilon^2)$.

We prove Theorem 3.1 by a reduction to the above communication problem, interpreting the one-way protocol as a sketching algorithm, as follows. Given the instance $(s_1, \ldots, s_{n/2}, i, t)$, define an n-vertex graph G that is a disjoint union of the graphs $\{G_j : j \in [\varepsilon^2 n/2]\}$, where each G_j is a bipartite graph, whose two sides, denoted $L(G_j)$ and $R(G_j)$, are of size

$$|L(G_j)| = |R(G_j)| = 1/\varepsilon^2.$$

The edges of G are determined by $s_1, \ldots, s_{n/2}$, where each string s_u is interpreted as a vector of indicators for the adjacency between vertex $u \in \cup_{j \in [\varepsilon^2 n/2]} L(G_j)$ and the respective $R(G_j)$.

Observe that Alice can compute G without any communication, as this graph is completely determined by her input. She then builds a sketch of this graph, that with probability $\geq 99/100$,

[1] The probability is over the two algorithms' common randomness; equivalently, the two algorithms have access to a public source of random bits

[2] Alice's input and Bob's input are *not* independent, but the marginal distribution of each one is uniform over its domain, namely, $\{0,1\}^{(n/2)\times(1/\varepsilon^2)}$ and $[n] \times \{0,1\}^{1/\varepsilon^2}$, respectively.

succeeds in simultaneously approximating all cut queries within factor $1 \pm \gamma\varepsilon$, where $\gamma > 0$ is a small constant to be determined later. This sketch is obtained from the theorem's assumption about m-bit sketches by standard amplification of the success probability from 3/4 to 0.99 (namely, repeating $r = O(1)$ times independently and answering any query with the median value of the r answers). Alice then sends this $O(m)$-bit sketch to Bob.

Bob then uses his input i to compute $j = j(i) \in [\varepsilon^2 n/2]$ such that the graph G_j contains vertex i (i.e., the vertex whose neighbors are determined by s_i). Bob also interprets his input string t as a vector of indicators determining a subset $T \subseteq R(G_j)$. By construction of G_j , the neighbor sets $N(v)$ of the vertices $v \in L(G_j) \setminus \{i\}$ are uniformly distributed, independently of T and of each other; in particular, each $|N(v) \cap T|$ has a Binomial distribution $B(\frac{1}{\varepsilon^2}, \frac{1}{4})$.

Lemma 3.3. *Using the $O(m)$-bit sketch he received from Alice, Bob can compute a "list" $B \subset L(G_j)$ of size $|B| = \frac{1}{2}|L(G_j)| = \frac{1}{2\varepsilon^2}$, and with probability at least 0.96, this list contains at least $\frac{4}{5}$-fraction of the vertices in the set*

$$L_{\text{high}} := \{v \in L(G_j) : \ |N(v) \cap T| \geq \tfrac{1}{4\varepsilon^2} + \tfrac{c}{\varepsilon}\}. \tag{2}$$

Moreover, Bob uses no information about his input i other than $j = j(i)$.

Before proving the lemma, let us show how to use it to decide about $\Delta(s_i, t)$ and derive the theorem. We will need the following simple claim, which we prove in Appendix A.

Claim 3.4. *With probability at least 0.98, the relative size of L_{high} is $\frac{|L_{\text{high}}|}{|L(G_j)|} \in [\frac{1}{2} - 10c, \frac{1}{2}]$.*

We assume henceforth that the events described in the above lemma and claim indeed occur, which happens with probability at least 0.94. Notice that $\Delta(s_i, t) = \deg(i) + |T| - 2|N(i) \cap T|$.

Now suppose that $\Delta(s_i, t) \leq \frac{1}{2\varepsilon^2} - \frac{c}{\varepsilon}$. Then $|N(i) \cap T| \geq \frac{1}{4\varepsilon^2} + \frac{c}{2\varepsilon}$, and because Bob's list B is independent of the vertex $i \in L(G_j)$, we have $\Pr[i \in B] \geq \frac{4}{5}|L_{\text{high}}|/|L_{\text{high}}| = \frac{4}{5}$.

Next, suppose that $\Delta(s_i, t) \geq \frac{1}{2\varepsilon^2} + \frac{c}{\varepsilon}$. Then $|N(i) \cap T| \leq \frac{1}{4\varepsilon^2} - \frac{c}{2\varepsilon}$, and using Claim 3.4,

$$\Pr[i \in B] \leq \frac{|B| - \frac{4}{5}|L_{\text{high}}|}{|L(G_j)|} \leq \frac{1}{4}.$$

Thus, Bob can decide between the two cases with error probability at most 1/4. Overall, it follows that Bob can solve the Gap Hamming problem for (s_i, t), with overall error probability at most $1/4 + 0.06 < 1/3$, as required to prove the theorem.

Proof of Lemma 3.3. We now show how Bob creates the "list" $B \subset L(G_j)$ of size $|B| = \frac{1}{2\varepsilon^2}$. Bob estimates the cut value for $S \cup T$ for every subset $S \subseteq L(G_j)$ of size exactly $\frac{1}{2\varepsilon^2}$. Observe that the cut value for a given S is

$$\delta(S \cup T) = \sum_{v \in S} \deg(v) + \sum_{u \in T} \deg(u) - 2 \sum_{v \in S} |N(v) \cap T|.$$

The cut value is bounded by the number of edges in G_j, which is at most $1/\varepsilon^4$, and since the data structure maintains all the cut values within factor $1 + \gamma\varepsilon$ for an arbitrarily small constant $\gamma > 0$, the additive error on each cut value is at most γ/ε^3. Further, we can assume Bob knows the exact degrees of all vertices (by adding them to the sketch, using $O(n \log \frac{1}{\varepsilon})$ bits), which he

can subtract off, and since scaling by $-1/2$ can only shrink the additive error, we can define the "normalized" cut value

$$n(S,T) := \sum_{v \in S} |N(v) \cap T|,$$

which Bob can estimate within additive error γ/ε^3. Bob's algorithm is to compute these estimates for all the values $n(S,T)$, and output a set S that maximizes his estimate for $n(S,T)$ as the desired list $B \subset L(G_j)$.

Let us now analyze the success probability of Bob's algorithm. For each vertex $v \in L(G_j)$, let $f(v) := |N(v) \cap T|$. Observe that each $f(v)$ has a Binomial distribution $B(\frac{1}{\varepsilon^2}, \frac{1}{4})$, and they are independent of each other. We will need the following bounds on the typical values of some order statistics when taking multiple samples from such a Binomial distribution. Recall that the *k-th order statistic* of a sample (multiset) $x_1, \ldots, x_m \in \mathbb{R}$ is the k-th smallest element in that sample. The following claim is proved in Appendix A.

Claim 3.5. *Let $\{X_j\}_{j=1,\ldots,m}$ be independent random variables with Binomial distribution $B(t, \frac{1}{4})$. Let $\alpha \in (0, \frac{1}{2})$ such that $(\frac{1}{2} + \alpha)m$ is integral, and both $t, m \geq 10/\alpha^2$. Then*

$$\Pr\Big[\textit{the } (\tfrac{1}{2} - \alpha)m \textit{ order statistic of } \{X_j\} \textit{ is} \leq \tfrac{1}{4}t - \tfrac{\alpha}{10}\sqrt{t}\Big] \geq 0.99, \textit{ and}$$
$$\Pr\Big[\textit{the } (\tfrac{1}{2} + \alpha)m \textit{ order statistic of } \{X_j\} \textit{ is} \geq \tfrac{1}{4}t + \tfrac{\alpha}{10}\sqrt{t}\Big] \geq 0.99.$$

Sort the vertices $v \in L(G_j)$ by their $f(v)$ value, i.e., denote them by $v_1, \ldots, v_{1/\varepsilon^2}$ such that $f(v_i) \leq f(v_{i+1})$. Applying the claim (for $\alpha = 0.05$ and $t, m = \frac{1}{\varepsilon^2}$), we see that with probability at least 0.98, the difference

$$f(v_{0.55/\varepsilon^2}) - f(v_{0.45/\varepsilon^2}) \geq 0.01/\varepsilon. \tag{3}$$

We assume henceforth this event indeed occurs. Let S^* include the $\frac{1}{2\varepsilon^2}$ vertices $v \in L(G_j)$ with largest $f(v)$, i.e., $S^* := \{v_j\}_{j > 0.5/\varepsilon^2}$, and let $S' \subset L(G_j)$ be any subset of the same size such that at least $\frac{1}{10}$-fraction of its vertices are not included in S^* (i.e., their order statistic in $L(G_j)$ is at most $\frac{1}{2\varepsilon^2}$). Then we can write

$$n(S^*,T) = \sum_{j \in S^*} f(v) = \sum_{j > 0.5/\varepsilon^2} f(v_j),$$
$$n(S',T) = \sum_{j \in S'} f(v) \leq \sum_{j > 0.6/\varepsilon^2} f(v_j) + \sum_{0.4/\varepsilon^2 < j \leq 0.5/\varepsilon^2} f(v_j).$$

Now subtract them

$$n(S',T) - n(S^*,T) = \sum_{0.5/\varepsilon^2 < j \leq 0.6/\varepsilon^2} f(v_j) - \sum_{0.4/\varepsilon^2 < j \leq 0.5/\varepsilon^2} f(v_j),$$

observe that elements in the normalized interval $(0.5, 0.55]$ dominate those in $(0.45, 0.5]$,

$$\geq \sum_{0.55/\varepsilon^2 < j \leq 0.6/\varepsilon^2} f(v_j) - \sum_{0.4/\varepsilon^2 < j \leq 0.45/\varepsilon^2} f(v_j)$$

and bound the remaining elements using (3),

$$\geq (0.05/\varepsilon^2)\big[f(v_{0.55/\varepsilon^2}) - f(v_{0.45/\varepsilon^2})\big] \geq 0.0005/\varepsilon^3.$$

Bob's estimate for each of the values $n(S^*, T)$ and $n(S', T)$ has additive error at most γ/ε^3, and therefore for suitable $\gamma = 10^{-4}$, the list B Bob computes cannot be this set S'. Thus, Bob's list B must contain at least 9/10-fraction of S^*, i.e., the $\frac{1}{2\varepsilon^2}$ vertices $v \in L(G_j)$ with highest $f(v)$.

Recall from Claim 3.4 that with probability at least 0.98, we have $\frac{1}{4\varepsilon^2} \leq |L_{\text{high}}| \leq \frac{1}{2\varepsilon^2}$, and assume henceforth this event occurs. Since S^* includes the $\frac{1}{2\varepsilon^2}$ vertices with highest f-value, it must contain all the vertices of L_{high}, i.e., $L_{\text{high}} \subseteq S^*$. We already argued that Bob's list B contains all but at most $\frac{1}{10}|S^*| = \frac{1}{20\varepsilon^2}$ vertices of S^*, and thus

$$\frac{|L_{\text{high}} \setminus B|}{|L_{\text{high}}|} \leq \frac{|S^* \setminus B|}{|L_{\text{high}}|} \leq \frac{\frac{1}{20\varepsilon^2}}{\frac{1}{4\varepsilon^2}} = \frac{1}{5}.$$

This bound holds with probability at least 0.96 (because of two events that we ignored, each having probability at most 0.02) and this proves the lemma. □

3.1 The communication lower bound

In this section we prove Theorem 3.2 (see Theorem 3.9 below), by considering distributional communication problems between two parties, Alice and Bob, as defined below. We restrict attention to the one-way model, in which Alice sends to Bob a single message M that is a randomized function of her input (using her private randomness), and Bob outputs the answer.

Distributional versions of Gap-Hamming Distance. Let $0 < c < 1$ be a parameter, considered to be a sufficiently small constant (Our analysis is asymptotic for ε tending to 0.) Alice's input is $S \in \{0,1\}^{\frac{1}{\varepsilon^2}}$, Bob's input is $T \in \{0,1\}^{\frac{1}{\varepsilon^2}}$, where the Hamming weights are $\text{wt}(S) = \text{wt}(T) = \frac{1}{2\varepsilon^2}$, and Bob needs to evaluate the partial function

$$f_c(S,T) = \begin{cases} 1 & \text{if } \Delta(S,T) \geq \frac{1}{2\varepsilon^2} + \frac{c}{\varepsilon}; \\ 0 & \text{if } \Delta(S,T) \leq \frac{1}{2\varepsilon^2} - \frac{c}{\varepsilon}. \end{cases}$$

The distribution μ we place on the inputs (S,T) is the following: S is chosen uniformly at random with $\text{wt}(S) = \frac{1}{2\varepsilon^2}$, and then with probability $\frac{1}{2}$, we choose T uniformly at random with $\text{wt}(T) = \frac{1}{2\varepsilon^2}$ subject to the constraint that $\Delta(S,T) \geq \frac{1}{2\varepsilon^2} + \frac{c}{\varepsilon}$, while with the remaining probability $\frac{1}{2}$, we choose T uniformly at random with $\text{wt}(T) = \frac{1}{2\varepsilon^2}$ subject to the constraint that $\Delta(S,T) \leq \frac{1}{2\varepsilon^2} - \frac{c}{\varepsilon}$. We say Alice's message $M = M(S)$ is δ-error for (f_c, μ) if Bob has a reconstruction function R for which

$$\Pr_{(S,T)\sim\mu,\ \text{private randomness}}[R(M,T) = f_c(S,T)] \geq 1 - \delta.$$

Now consider a related but different distributional problem. Alice and Bob have $S, T \in \{0,1\}^{\frac{1}{\varepsilon^2}}$, respectively, each of Hamming weight exactly $\frac{1}{2\varepsilon^2}$, and Bob needs to evaluate the function

$$g(S,T) = \begin{cases} 1 & \text{if } \Delta(S,T) > \frac{1}{2\varepsilon^2}; \\ 0 & \text{if } \Delta(S,T) \leq \frac{1}{2\varepsilon^2}. \end{cases}$$

We place the following distribution ζ on the inputs (S,T): S and T are chosen independently and uniformly at random among all vectors with Hamming weight exactly $\frac{1}{2\varepsilon^2}$. We say a message M is δ-error for (g,ζ) if Bob has a reconstruction function R for which

$$\Pr_{(S,T)\sim\zeta,\text{ private randomness}}[R(M,T)=g(S,T)]\geq 1-\delta.$$

Let $I(S;M)=H(S)-H(S|M)$ be the mutual information between S and M, where H is the entropy function. Define $\mathrm{IC}_{\mu,\delta}(f_c):=\min_{\delta\text{-error } M \text{ for } (f_c,\mu)} I(S;M)$ and $\mathrm{IC}_{\zeta,\delta}(g):=\min_{\delta\text{-error } M \text{ for } (g,\zeta)} I(S;M)$.

Lemma 3.6. *For all $\delta>0$, $\mathrm{IC}_{\mu,\delta}(f_c)\geq \mathrm{IC}_{\zeta,\delta+O(c)}(g)$.*

Proof. It suffices to show that if M is δ-error for (f_c,μ), then M is $(\delta+O(c))$-error for (g,ζ). Since M is δ-error for (f_c,μ), Bob has a reconstruction function R for which

$$\Pr_{(S,T)\sim\mu,\text{ private randomness}}[R(M,T)=f_c(S,T)]\geq 1-\delta.$$

Now consider $\Pr_{(S,T)\sim\zeta,\text{ private randomness}}[R(M,T)=g(S,T)]$. Observe that whenever (S,T) lies in the support of μ, if $R(M,T)=f_c(S,T)$, then $R(M,T)=g(S,T)$. The probability that (S,T) lies in the support of μ is $1-O(c)$, by standard anti-concentration arguments of the Binomial distribution (or the Berry-Esseen Theorem), and conditioned on this event we have that (S,T) is distributed according to μ. Hence, $\Pr_{(S,T)\sim\zeta,\text{ private randomness}}[R(M,T)=g(S,T)]\geq[1-O(c)][1-\delta]\geq 1-O(c)-\delta$. □

We now lower bound $\mathrm{IC}_{\zeta,\delta}(g)$.

Lemma 3.7. *For $\delta_0>0$ a sufficiently small constant, $IC_{\zeta,\delta_0}(g)=\Omega\left(\frac{1}{\varepsilon^2}\right)$.*

Proof. We use the following lower bound of Braverman, Garg, Pankratov and Weinstein [BGPW13] for the following $h_c(S,T)$ problem. Like before, Alice has $S\in\{0,1\}^{\frac{1}{\varepsilon^2}}$, Bob has $T\in\{0,1\}^{\frac{1}{\varepsilon^2}}$, and needs to evaluate the partial function

$$h_c(S,T)=\begin{cases}1 & \text{if } \Delta(S,T)\geq\frac{1}{2\varepsilon^2}+\frac{c}{\varepsilon};\\ 0 & \text{if } \Delta(S,T)\leq\frac{1}{2\varepsilon^2}-\frac{c}{\varepsilon}.\end{cases}$$

However, now $\mathrm{wt}(S)$ and $\mathrm{wt}(T)$ may be arbitrary. Moreover, S and T are chosen independently and uniformly at random from $\{0,1\}^{\frac{1}{\varepsilon^2}}$. Denote this by $(S,T)\sim\eta$. Now it may be the case that $\left|\Delta(S,T)-\frac{1}{2\varepsilon^2}\right|<\frac{c}{\varepsilon}$, in which case Bob's output is allowed to be arbitrary. A message M is δ-error for (h_c,η) if Bob has a reconstruction function R for which

$$\Pr_{(S,T)\sim\eta,\text{ private randomness}}\left[(R(M,T)=h_c(S,T))\wedge\left(\left|\Delta(S,T)-\frac{1}{2\varepsilon^2}\right|\geq\frac{c}{\varepsilon}\right)\right]\geq 1-\delta.$$

It was proved in [BGPW13] that for a sufficiently small constant $\delta>0$,

$$\mathrm{IC}_{\eta,\delta}(h_1):=\min_{\delta\text{-error } M \text{ for } (h_1,\eta)} I(S;M)\geq\frac{C}{\varepsilon^2},$$

for an absolute constant $C>0$. We show how to apply this result to prove the lemma.

An immediate corollary of this result is that $\mathrm{IC}_{\eta,\delta}(g) := \min_{\delta\text{-error } M \text{ for } (g,\,\eta)} I(S;M) \geq \frac{C}{\varepsilon^2}$. Indeed, if M is δ-error for (g,η), then it is also δ-error for (h_1,η).

Now let M be a δ-error protocol for (g,ζ). Consider the following randomized protocol M' for g with inputs distributed according to η. Given S, Alice computes $s = \mathrm{wt}(S)$. If $s < \frac{1}{2\varepsilon^2}$, Alice randomly chooses $\frac{1}{2\varepsilon^2} - s$ coordinates in S that are equal to 0 and replaces them with a 1, otherwise she randomly chooses $s - \frac{1}{2\varepsilon^2}$ coordinates in S that are equal to 1 and replaces them with a 0. Let S' be the resulting vector. Alice sends $M(S')$ to Bob, i.e., $M'(S) := M(S')$. Given the message $M(S')$ and his input T, Bob first computes $t = \mathrm{wt}(T)$. If $t < \frac{1}{2\varepsilon^2}$, Bob randomly chooses $\frac{1}{2\varepsilon^2} - t$ coordinates in T which are equal to 0 and replaces them with a 1, otherwise he randomly chooses $t - \frac{1}{2\varepsilon^2}$ coordinates in T which are equal to 1 and replaces them with a 0. Let T' be the resulting vector. Suppose R is such that $\Pr_{(S',T')\sim\zeta,\text{ private randomness}}[R(M(S'),T') = g(S',T')] \geq 1-\delta$. Bob outputs $R(M(S'),T')$.

We now lower bound $\Pr[g(S',T') = g(S,T)]$, where the probability is over $(S,T)\sim\eta$ and the random choices of Alice and Bob for creating S',T' from S,T, respectively. First, the number of coordinates changed by Alice or Bob is $r = \Theta(1/\varepsilon)$ with arbitrarily large constant probability. Since S and T are independent and uniformly random, after performing this change, the Hamming distance on these r coordinates is $\frac{r}{2} \pm O(\sqrt{r})$ with arbitrarily large constant probability. Finally, $\left|\Delta(S',T') - \frac{1}{2\varepsilon^2}\right| = \omega(\sqrt{r})$ with arbitrarily large constant probability. Hence, with arbitrarily large constant proabbiltiy, $g(S',T') = g(S,T)$. It follows that $\Pr[g(S',T') = g(S,T)] \geq 1-\gamma$ for an arbitrarily small constant $\gamma > 0$, and therefore if R' describes the above reconstruction procedure of Bob, then $\Pr_{(S,T)\sim\eta,\text{ private randomness}}[R'(M'(S),T) = g(S,T)] \geq 1-\gamma-\delta$.

Hence, M' is a $(\delta+\gamma)$-error protocol for (g,η). We now bound $I(M';S)$ in terms of $I(M;S')$. Let J be an indicator random variable for the event $\mathrm{wt}(S) \in \left[\frac{1}{2\varepsilon^2} - \frac{1}{\varepsilon^{3/2}}, \frac{1}{2\varepsilon^2} + \frac{1}{\varepsilon^{3/2}}\right]$. Then $\Pr[J = 1] = 1 - o(1)$, where $o(1) \to 0$ as $\varepsilon \to 0$. Since conditioning on a random variable Z can change the mutual information by at most $H(Z)$, we have

$$I(M';S) \leq I(M';S \mid J) + H(J) \leq I(M';S \mid J = 1) + 1. \tag{4}$$

S is a probabilistic function of S', which if $J = 1$, is obtained by changing at most $1/\varepsilon^{3/2}$ randomly chosen coordinates $A_1,\ldots,A_{1/\varepsilon^{3/2}}$ of S' from 0 to 1 or from 1 to 0. By the data processing inequality and the chain rule for mutual information,

$$\begin{aligned} I(M';S \mid J=1) &\leq I(M';S',A_1,\ldots,A_{1/\varepsilon^{3/2}} \mid J=1) \\ &= I(M';S' \mid J=1) + \sum_{\ell=1}^{1/\varepsilon^{3/2}} I(M';A_\ell \mid J=1,A_1,\ldots,A_{\ell-1}) \\ &\leq I(M';S' \mid J=1) + O\left(\frac{\log(1/\varepsilon)}{\varepsilon^{3/2}}\right). \end{aligned} \tag{5}$$

Observe that the joint distribution of $M'(S')$ and S' is independent of J, and moreover is equal to the joint distribution of $M(S')$ and $S'\sim\zeta$. We can take M to be a δ-error protocol for (g,ζ) for which $I(M(S');S') = \mathrm{IC}_{\zeta,\delta}(g)$. Combining this with (4) and (5), $I(M';S) \leq \mathrm{IC}_{\zeta,\delta}(g) + O\left(\frac{\log(1/\varepsilon)}{\varepsilon^{3/2}}\right)$. Now since M' is a $(\delta+\gamma)$-error protocol for (g,η), we have $I(M';S) \geq \mathrm{IC}_{\eta,\delta+\gamma}(g) \geq \frac{C}{\varepsilon^2}$, provided δ and γ are sufficiently small constants. It follows that $\mathrm{IC}_{\zeta,\delta}(G) \geq \frac{C}{\varepsilon^2} - O\left(\frac{\log(1/\varepsilon)}{\varepsilon^{3/2}}\right) \geq \frac{C}{2\varepsilon^2}$, as desired. □

Corollary 3.8. *For sufficiently small constants $\delta, c > 0$, $\mathrm{IC}_{\mu,\delta}(f_c) = \Omega(1/\varepsilon^2)$.*

Proof. This follows by combining Lemmas 3.6 and 3.7. □

n-fold version of Gap-Hamming Distance. We now consider the n-fold problem in which Alice is given n strings $S_1, \ldots, S_n \in \{0,1\}^{1/\varepsilon^2}$, and Bob has an index $I \in [n]$ together with one string $T \in \{0,1\}^{1/\varepsilon^2}$. Here $(S_I, T) \sim \zeta$, while S_j for $j \neq I$, are chosen independently and uniformly at random from all Hamming weight $\frac{1}{2\varepsilon^2}$ vectors. Thus the joint distribution of $S_1, \ldots, S_n$ is n i.i.d. strings drawn uniformly from $\{0,1\}^{1/\varepsilon^2}$ subject to each of their Hamming weights being $\frac{1}{2\varepsilon^2}$. Here I is drawn independently and uniformly from $[n]$. We let ν denote the resulting input distribution.

We consider the one-way two-party model in which Alice sends a single, possibly randomized message M of her inputs $S_1, \ldots, S_n$, and Bob needs to evaluate $h(S_1, \ldots, S_n, T) = f_c(S_I, T)$. We say M is δ-error for (h, ν) if Bob has a reconstruction function R for which

$$\Pr_{\text{inputs}\sim\nu,\text{ private randomness}}\left[(R(M,T,I) = f_c(S_I,T)) \wedge \left(\left|\Delta(S_I,T) - \frac{1}{2\varepsilon^2}\right| \geq \frac{c}{\varepsilon}\right)\right] \geq 1 - \delta.$$

Let $\mathrm{IC}_{\nu,\delta}(h) := \min_{\delta\text{-error } M \text{ for } (h,\nu)} I(S_1, \ldots, S_n; M)$.

Theorem 3.9. *For a sufficiently small constant $\delta > 0$, $\mathrm{IC}_{\nu,\delta}(h) = \Omega(n/\varepsilon^2)$. In particular, the distributional one-way communication complexity of h under input distribution ν is $\Omega(n/\varepsilon^2)$.*

Proof. Say an index $i \in [n]$ is *good* if

$$\Pr_{\text{inputs}\sim\nu,\text{ private randomness}}\left[(R(M,T,I) = f_c(S_I,T)) \wedge \left(\left|\Delta(S_I,T) - \frac{c}{2\varepsilon^2}\right| \geq \frac{1}{\varepsilon}\right) \mid I = i\right] \geq 1 - 2\delta.$$

By a union bound, there are at least $n/2$ good $i \in [n]$. By the chain rule for mutual information and using that the S_i are independent and conditioning does not increase entropy,

$$I(M; S_1, \ldots, S_n) \geq \sum_{i=1}^{n} I(M; S_i) \geq \sum_{\text{good } i} I(M; S_i).$$

We claim that for each good i, $I(M; S_i) \geq \mathrm{IC}_{\mu,2\delta}(f_c)$. Indeed, consider the following protocol M_i for f_c under distribution μ. Alice, given her input S for f_c, uses her private randomness to sample S_j for all $j \neq i$ independently and uniformly at random from $\{0,1\}^{1/\varepsilon^2}$ subject to each of their Hamming weights being $\frac{1}{2\varepsilon^2}$. Bob sets $I = i$ and uses his input T for f_c as his input for h. Since i is good, it follows that M_i is 2δ-error for (f_c, ζ). Hence $I(M; S_i) \geq \mathrm{IC}_{\mu,2\delta}(f_c)$, which by Corollary 3.8, is $\Omega(1/\varepsilon^2)$ provided $\delta > 0$ is a sufficiently small constant. Hence, $\mathrm{IC}_{\nu,\delta} = \Omega(n/\varepsilon^2)$. Since $\mathrm{IC}_{\nu,\delta}(h) \leq I(M; S_1, \ldots, S_n)$ for each δ-error M for (h, ν), and $I(M; S_1, \ldots, S_n) \leq H(M)$ which is less than the length of M, the communication complexity lower bound follows. □

References

[AG09] K. J. Ahn and S. Guha. Graph sparsification in the semi-streaming model. In *ICALP (2)*, pages 328–338, 2009.

[AGM12] K. J. Ahn, S. Guha, and A. McGregor. Graph sketches: sparsification, spanners, and subgraphs. In *PODS*, pages 5–14, 2012.

[Alo97] N. Alon. On the edge-expansion of graphs. *Comb. Probab. Comput.*, 6(2):145–152, June 1997. doi:10.1017/S096354839700299X.

[ARV09] S. Arora, S. Rao, and U. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2):1–37, 2009. doi:http://doi.acm.org/10.1145/1502793.1502794.

[BGPW13] M. Braverman, A. Garg, D. Pankratov, and O. Weinstein. Information lower bounds via self-reducibility. In *Computer Science Theory and Applications*, volume 7913 of *Lecture Notes in Computer Science*, pages 183–194. Springer, 2013. doi:10.1007/978-3-642-38536-0_16.

[BK96] A. A. Benczúr and D. R. Karger. Approximating s-t minimum cuts in $\tilde{O}(n^2)$ time. In *28th Annual ACM Symposium on Theory of Computing*, pages 47–55. ACM, 1996. doi:10.1145/237814.237827.

[BK02] A. A. Benczúr and D. R. Karger. Randomized approximation schemes for cuts and flows in capacitated graphs. *CoRR*, cs.DS/0207078, 2002. Available from: http://arxiv.org/abs/cs.DS/0207078.

[BSS12] J. D. Batson, D. A. Spielman, and N. Srivastava. Twice-ramanujan sparsifiers. *SIAM J. Comput.*, 41(6):1704–1721, 2012. doi:10.1137/090772873.

[FHHP11] W. S. Fung, R. Hariharan, N. J. Harvey, and D. Panigrahi. A general framework for graph sparsification. In *43rd Annual ACM Symposium on Theory of Computing*, pages 71–80. ACM, 2011. doi:10.1145/1993636.1993647.

[GKK10] A. Goel, M. Kapralov, and S. Khanna. Graph sparsification via refinement sampling. *arXiv preprint arXiv:1004.4915*, 2010.

[GRV09] N. Goyal, L. Rademacher, and S. Vempala. Expanders via random spanning trees. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 576–585, 2009.

[Kar00] D. R. Karger. Minimum cuts in near-linear time. *Journal of the ACM (JACM)*, 47(1):46–76, 2000.

[KL02] D. R. Karger and M. S. Levine. Random sampling in residual graphs. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 63–66, 2002.

[KL13] J. A. Kelner and A. Levin. Spectral sparsification in the semi-streaming setting. *Theory Comput. Syst.*, 53(2):243–262, 2013.

[KMP10] I. Koutis, G. L. Miller, and R. Peng. Approaching optimality for solving SDD linear systems. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 235–244, 2010.

[KMP11] I. Koutis, G. L. Miller, and R. Peng. A nearly-$m \log n$ time solver for SDD linear systems. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 590–598, 2011.

[KP12] M. Kapralov and R. Panigrahy. Spectral sparsification via random spanners. In *3rd Innovations in Theoretical Computer Science Conference*, pages 393–398. ACM, 2012. doi:10.1145/2090236.2090267.

[LPS88] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[Mad10] A. Madry. Fast approximation algorithms for cut-based problems in undirected graphs. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 245–254. IEEE, 2010.

[MSS13] A. Marcus, D. Spielman, and N. Srivastava. Interlacing families I: Bipartite Ramanujan graphs of all degrees. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, 2013. To appear.

[Nil91] A. Nilli. On the second eigenvalue of a graph. *Discrete Math*, 91:207–210, 1991. doi:10.1016/0012-365X(91)90112-F.

[She09] J. Sherman. Breaking the multicommodity flow barrier for $O(\sqrt{\log n})$-approximations to sparsest cut. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 363–372, 2009.

[SS11] D. A. Spielman and N. Srivastava. Graph sparsification by effective resistances. *SIAM J. Comput.*, 40(6):1913–1926, December 2011. doi:10.1137/080734029.

[ST04] D. A. Spielman and S.-H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *36th Annual ACM Symposium on Theory of Computing*, pages 81–90. ACM, 2004. doi:10.1145/1007352.1007372.

[ST11] D. A. Spielman and S.-H. Teng. Spectral sparsification of graphs. *SIAM J. Comput.*, 40(4):981–1025, July 2011. doi:10.1137/08074489X.

[Vit85] J. S. Vitter. Random sampling with a reservoir. *ACM Trans. Math. Softw.*, 11(1):37–57, March 1985. doi:10.1145/3147.3165.

A Additional Proofs

Proof of Claim 3.4. By basic properties of the Binomial distribution (or the Berry-Esseen Theorem), there are absolute constants $\frac{1}{5} \le K_1 \le K_2 \le 5$ such that for each vertex $v \in L(G_j)$,

$$\Pr\left[v \in L_{\text{high}}\right] = \Pr\left[|N(v) \cap T| \ge \tfrac{1}{4\varepsilon^2} + \tfrac{c}{\varepsilon}\right] \in [\tfrac{1}{2} - K_2 c, \tfrac{1}{2} - K_1 c].$$

Denoting $Z := |L_{\text{high}}|$, we have by Hoeffding's inequality,

$$\Pr\left[|Z - \mathbb{E}[Z]| > \tfrac{K_1 c}{\varepsilon^2}\right] \le 2e^{-\frac{1}{2}(K_1 c)^2 (1/\varepsilon^2)} \le 0.02.$$

Thus, with probability at least 0.98, we have both bounds

$$Z \le \mathbb{E}[Z] + \tfrac{K_1 c}{\varepsilon^2} \le \tfrac{1}{2\varepsilon^2}, \text{ and}$$
$$Z \ge \mathbb{E}[Z] - \tfrac{K_1 c}{\varepsilon^2} \le \tfrac{1}{2\varepsilon^2} \ge (\tfrac{1}{2} - 2K_2 c)\tfrac{1}{\varepsilon^2} \ge \tfrac{1}{4\varepsilon^2}.$$

□

Proof of Claim 3.5. The $(\frac{1}{2} - \alpha)m$ order statistic of $\{X_j\}$ is smaller or equal to $T := \frac{1}{4}t - \frac{\alpha}{10}\sqrt{t}$ if and only if at least $(\frac{1}{2} - \alpha)m$ elements are smaller or equal to T, which can be written as $\sum_{j=1}^{m} \mathbb{1}_{\{X_j \le T\}} \ge (\frac{1}{2} - \alpha)m$.

Now fix $j \in \{1, \ldots, t\}$. Then

$$\Pr[X_j \le T] = \Pr[X_j \le \tfrac{1}{4}t] \cdot \Pr[X_j \le T \mid X_j \le \tfrac{1}{4}t], \tag{6}$$

and by the Binomial distribution's relationship between mean and median, $\Pr[X_j \le \frac{1}{4}t] \ge \frac{1}{2}$. Elementary but tedious calculations (or the Berry-Esseen Theorem) show there is an absolute constant $K \in (0,5)$ such that

$$\Pr\left[\tfrac{1}{4}t - \tfrac{\alpha}{10}\sqrt{t} < X_j \le \tfrac{1}{4}t\right] \le K\tfrac{\alpha}{10} \cdot \Pr\left[X_j \le \tfrac{1}{4}t\right],$$

and plugging into (6), we obtain $\Pr[X_j \le T] \ge \frac{1}{2}(1 - K\frac{\alpha}{10}) \ge \frac{1}{2} - \frac{1}{2}\alpha$.

Now bound the expectaton by $\mathbb{E}[\sum_{j=1}^{m} \mathbb{1}_{\{X_j \le T\}}] \ge (\frac{1}{2} - \frac{1}{2}\alpha)m$, and apply Hoeffding's inequality,

$$\Pr\Big[\sum_j \mathbb{1}_{\{X_j \le T\}} < (\tfrac{1}{2} - \alpha)m\Big] \le e^{-\frac{1}{2}(\frac{1}{2}\alpha)^2 m} = e^{-\alpha^2 m/8} \le 0.01,$$

where the last inequality follows since $\alpha^2 m$ is sufficiently large. □

B Sketching Algorithm for General Edge Weights

We now show the upper bound (Theorem 2.1) for general edge weights. Assume there is a sketching algorithm, which we shall call the "basic sketch", for the case where all edge weights are in a polynomial range, say for concreteness $[1, n^5]$ (which by scaling is equivalent to the range $[b, n^5 b]$ for any $b > 0$), which uses space $\tilde{O}(n/\varepsilon)$. We may assume the success probability of this sketch is at least $1 - 1/n^8$, e.g., by using standard amplification using $O(\log n)$ repetitions, thereby increasing the sketch size by at most $O(\log n)$ factor. Throughout, we measure the memory usage in machine words, and assume that a machine word can accomodate the weight of an edge and also at least $2\log n$ bits. As before, we may assume $\varepsilon > 1/n$, as otherwise the theorem is trivial.

B.1 Construction

Sketching Algorithm. The sketch has two components. The first component is essentially a maximum-weight spanning tree T computed using Kruskal's algorithm. Specifically, start with T as an empty graph on vertex set V (so every vertex forms a connected component of size one), and then go over the edges $e \in E$ in some order π of decreasing weight, each time adding e to the current T if this would not close a cycle in T (i.e., the endpoints of e are currently in different connected components). If the edge weights are all distinct, then π is unique; otherwise, fix some π by breaking ties arbitrarily (and use the same π later, for sake of consistency). At the end, T is a spanning tree and has $n-1$ edges. Our sketch stores the list of edges that form T, denoted $e_1, e_2, \ldots, e_{n-1} \in E$, sorted in order of insertion (which is also their ordering according to π).

The second component is computed by iterating over $e_1, \ldots, e_{n-1}$; in iteration j, we consider the graph obtained from G in three steps: (i) remove all edges $e \in E$ of weight $w(e) < w(e_j)/n^3$; (ii) change all edges $e \in E$ of weight $w(e) \ge n^2 \cdot w(e_j)$ to have infinite weight; and (iii) contract all edges of infinite weight, keeping parallel edges (self-loops may be removed as they will have no effect anyway).[3] We denote the graph obtained after step (ii) by G_j, and the one obtained after step (iii) by G'_j. Notice that from the perspective of cuts, G_j and G'_j are equivalent, and we thus refer to G'_j as the *reduced form* of G_j. Observe that G'_j has at most n vertices, and its edge weights lie in

[3] We manage vertex names in a systematic manner, e.g., a merging of vertices u, v keeps the lexicographically smaller name.

the range $[n^{-3} \cdot w(e_j), n^2 \cdot w(e_j)]$, so in principle, we can apply the assumed sketching algorithm (that works for edge weight in a polynomial range) on G'_j. We do so (apply the assumed sketching algorithm) and store its result in our sketch, but with two twists: First, we apply the assumed sketching algorithm separately on every connected component of G'_j of size at least 2 rather than on the entire G'_j. Second, we do it only if there is no earlier iteration $k < j$ with $w(e_k)/w(e_j) < 2$ for which we already sketched and stored G'_k. For instance, if $w(e_{j-1}) = w(e_j)$, then G'_{j-1} and G'_j are identical and G'_j will not be stored (because of $k = j - 1$ or an even smaller k).

Estimation algorithm. Given a query subset $S \subset V$, find the smallest $j \in [n-1]$ such that e_j crosses the cut $(S, \bar{S})$; such j exists because $\{e_1, \ldots, e_{n-1}\}$ forms a spanning tree. We further show in Lemma B.1 below that e_j is a heaviest edge in this cut, hence $w(S, \bar{S})/w(e_j) \in [1, n^2]$. Now find the largest $k \leq j$ for which we sketched and stored G'_k; by construction $w(e_k)/w(e_j) \in [1, 2)$. Lemma B.2 below proves that the cut values in G and in G_k are almost the same. Next, compute the connected components of the graph $(V, \{e_1, \ldots, e_k\})$, and observe they must be exactly the same as the connected components of G_k. Obviously, the value of the cut $(S, \bar{S})$ in G_k is just the sum, over all connected components $V' \subset V$ in G_k, of the contribution to the cut from edges inside that component, namely $w(S \cap V', \bar{S} \cap V')$. Recall that G'_k has essentially the same cuts as G_k and we can thus estimate each such term $w(S \cap V', \bar{S} \cap V')$ using the sketch we prepared for G'_k (more precisely, using the sketch of the respective component V' of G'_k, unless $|V'| = 1$ in which case that term is trivially 0). To this end, we need to find out which vertices of G_k were merged together to form G'_k, which can be done using $e_1, \ldots, e_{n-1}$ as follows. Find the largest k^* such that $w(e_{k^*}) \geq n^2 \cdot w(e_k)$, and compute the connected components of the graph $(V, \{e_1, \ldots, e_{k^*}\})$. Lemma B.3 below proves that these connected components (or more precisely the partition of V they induce) are exactly the subsets of vertices that are merged in G_k to create G'_k. Now that know the vertex correspondence between G_k and G'_k, we estimate the cut value $w(S \cap V', \bar{S} \cap V')$ by simply using the estimate for the corresponding cut value in G'_k, where the latter is obtained using the basic sketch prepared for G'_k.

B.2 Accuracy Guarantee

The accuracy of the estimation algorithm follows from the above discussion, which uses the three lemmas below, together with a union bound over the events of an error in any of the basic estimates used along the way, the number of which is $O(n)$, because they correspond to disjoint subsets of V. (The union bound is applicable because these basic sketch are queried in a non-adaptive manner, or alternatively, because we make at most one query to every basic sketch that is constructed independetly of the others.)

Lemma B.1. *Fix $S \subset V$ and let $e' \in E$ be the first edge, according to the ordering π, that crosses the cut $(S, \bar{S})$. Then this e' is the first edge in the sequence $e_1, \ldots, e_{n-1}$ that crosses the cut $(S, \bar{S})$.*

Proof. Let $e' \in E$ be the first edge, according to the ordering π, that crosses the cut $(S, \bar{S})$. Clearly, e' is a heaviest edge in this cut. Now observe that in the construction of T (i.e., $e_1, \ldots, e_{n-1}$), when e' is considered, T has no edges between S and $\bar{S}$, hence the endpoints of e lie in different connected components, and e' must be added to T. □

Lemma B.2. *Consider a query $S \subset V$ and let $k \in [n-1]$ be the value computed in the estimation algorithm. Then the ratio between the value of $w(S, \bar{S})$ in the graph G_k and that in the graph G is*

in the range $[1-\frac{1}{n},1] \subset [1-\varepsilon,1]$*, formally*

$$1-\tfrac{1}{n} \leq \frac{w_{G_k}(S,\bar{S})}{w_G(S,\bar{S})} \leq 1.$$

Proof. The edges in G_k are obtained from the edges of G, by either (1) removing edges e whose weight is $w(e) < w(e_k)/n^3$; or (2) changing edges e with $w(e) \geq n^2 \cdot w(e_k)$ to have infinite weight. The first case can only decrease any cut value, while the second case can only increase any cut value.

Recall that the estimation algorithm finds j such that $w(S,\bar{S})/w(e_j) \in [1,n^2]$, and then finds $k \leq j$, which we said always satisfies $w(e_k)/w(e_j) \in [1,2)$. Thus, $w(S,\bar{S})/w(e_k) \in (\frac{1}{2},n^2]$. So one direction of the desired inequality follows by observing that edges in G that fall into case (1) have the total weight at most

$$\tbinom{n}{2} w(e_k)/n^3 \leq \tfrac{2}{n} w(e_k) \leq \tfrac{1}{n} w(S,\bar{S}).$$

The other direction follows by observing that edges e that fall into case (2) have (in G) weight $w(e) > n^2 \cdot w(e_k) \geq w(S,\bar{S})$, and therefore do not belong to the cut $(S,\bar{S})$. □

Lemma B.3. *Fix $w^* > 0$, let $E^* := \{e \in E : w(E) \geq w^*\}$, and find the largest $i^* \in [n-1]$ such that $w(e_{i^*}) > w^*$. Then the graphs (V,E^*) and $(V,\{e_1,\ldots,e_{i^*}\})$ have exactly the same connected components (in terms of the partition they induce of V).*

Proof. It is easy to see that executing our construction of T above on the set E^*, gives the exact same result as executing it for E but stopping once we reach edges of weight smaller than w^*. The latter results with the edges $e_1,\ldots,e_{i^*}$, while the former is clearly an execution of Kruskal's algorithm, i.e. computes a maximum weight forest in E^*. □

B.3 Size Analysis

Lemma B.4. *The total size of the sketch is at most $\tilde{O}(n/\varepsilon \log\log W)$, where we assume all non-zero edge weights are in the range $[1,W]$.*

Proof. The first component of the sketch is just a list of $n-1$ edges with their edge weights, hence its size is $O(n\log(\log W/\varepsilon))$ (we can store a $1+\varepsilon/2$ approximation to each weigth using space $\log(\log W/\varepsilon)$).

The second component of the sketch has $n-1$ parts, one for each G'_j where $j \in [n-1]$. For some of these j values, we compute and store the basic sketch for every connected components of G'_j that is of size at least 2. Denoting by n_j the number of vertices in G'_j, and by m_j the number of connected components in G'_j, the storage requirement for each G'_j is at most $\tilde{O}(\frac{n_j-m_j}{\varepsilon})$, because each connected component of size $s \geq 2$ requires storage $\tilde{O}(\frac{s}{\varepsilon}) \leq \tilde{O}(\frac{s-1}{\varepsilon})$, and these sizes (the different s values) add up to at most n.

Denote the values of j for which we do store a basic sketch for G'_j by $j_1 < j_2 < \cdots < j_p$, where by construction $w(e_{j_i})/w(e_{j_{i+1}}) \geq 2$. Summing over these values of j, the second component's storage requirement is at most

$$\tilde{O}(\tfrac{1}{\varepsilon} \sum_{i\in[p]} (n_{j_i} - m_{j_i})). \tag{7}$$

To ease notation, let $M := 5\log_2 n$, and consider the graphs G'_{j_i} and $G'_{j_{i+M}}$ for some $[i \in p-M-1]$. Observe that every edge in G'_{j_i} has weight at least $w(e_{j_i})/n^3 \geq 2^M \cdot w(e_{j_{i+M}})/n^3 = n^2 \cdot w(e_{j_{i+M}})$ (because edges of smaller weight are removed); thus, in $G_{j_{i+M}}$, these same edges have infinite weight, and then to create the reduced form $G_{j_{i+M}}$, these edges are contracted. It follows from this observation that every connected component in G'_{j_i} becomes in $G_{j_{i+M}}$ a single vertex, hence $n_{j_{i+M}} \leq m_{j_i}$ (we do not obtain equality since additional contractions may occur). Using this last inequality, for every $i^* \in [M]$, we can bound the following by a telescopic sum

$$\sum_{i=i^*,i^*+M,i^*+2M,\ldots} (n_{j_i} - m_{j_i}) \leq \sum_{i=i^*,i^*+M,i^*+2M,\ldots} (n_{j_i} - n_{j_{i+M}}) \leq n_{j_{i^*}} \leq n,$$

and therefore

$$\sum_{i\in[p]} (n_{j_i} - m_{j_i})) \leq \sum_{i^*\in[M]} \sum_{i=i^*,i^*+M,i^*+2M,\ldots} (n_{j_i} - m_{j_i}) \leq M \cdot n.$$

Plugging this last inequality into (7), we obtain that the second component's storage requirement is at most $M \cdot \tilde{O}(n/\varepsilon)$, which is still bounded by $\tilde{O}(n/\varepsilon)$. □

C Application: Stream Algorithm for Approximating Minimum Cut

We show an application of our sketch from Theorem 2.1 to approximating the (global) minimum cut of an input graph G in the data-stream model (i.e., as a list of edges added one by one) using two passes and space $\tilde{O}(n/\varepsilon)$. Previous results have larger space requirement $\tilde{O}(n/\varepsilon^2)$ but need only one pass [AG09, KL13, AGM12, GKK10]. We conjecture that our algorithm may be implemented in one pass as well.

Theorem C.1. *For every $\varepsilon > 0$ there is a two-pass streaming algorithm that reads an n-vertex graph G with polynomially bounded edge weights using space $\tilde{O}(n/\varepsilon)$, and outputs (at the end of the second pass) a subset $S \subset V$ whose cut capacity $w(S, \bar{S})$ is within factor $1+\varepsilon$ of the minimum cut of G, with probability at least $9/10$.*

Proof. Our minimum-cut algorithm starts, as described below, by implementing the sketch from Theorem 2.1 in two passes over the graph G. But for the purpose of amplifying its success probability to $1 - 1/n^4$, we shall actually constuct in parallel $O(\log n)$ independent copies of the sketch described below, and call the result a "super-sketch".

In the first pass, the algorithm computes a standard cut sparsifier H^* with approximation $1 + 1/\log n$, using any of the known streaming algorithms [AG09, KL13, AGM12, GKK10]. At the same time, the algorithm follows the proof of Theorem 2.1, and for every cut value $c \in \tilde{C}$, it samples every edge e with some prescribed probability p_e (after rescaling the edge weights and discarding heavy edges), and filters the resulting stream to create the classes L_i for $i \in [l]$. Each of these classes is a stream of edges, which we may view as a graph denoted $G_{c,i}$, and build for it a standard factor 2 cut sparsifier, denoted $T_{c,i}$ (using again known streaming algorithms). Each sparsifier $T_{c,i}$ has $\tilde{O}(n)$ edges, and thus all them together use a total of $\tilde{O}(n)$ space.

In the second pass, we complete computing a sketch as described in Theorem 2.1. Specifically, for every cut value $c \in \tilde{C}$, and class $i \in [l]$, we need to "break" the graph $G_{c,i}$ along its sparse cuts,

which we can perform using the sparsifier $T_{c,i}$ as follows. Using known algorithms [ARV09, She09] that achieve ρ-approximation for $\rho = O(\sqrt{\log n})$, we recursively partition the graph whenever that algorithm finds a sparse cut T with value $\frac{|E(T,\bar{T})|}{|T|} \leq \rho/(2\varepsilon)$. When this recursion stops, we are guaranteed that in each part P of $G_{c,i}$ the sparse cut value is greater than ε (as we take into account the ρ-approximation and the factor 2 error of the sparsifier $T_{c,i}$). We call the resulting partition $\mathcal{P}_{c,i}$, and observe that our recursive process breaks up sparse cuts with value in $G_{i,c}$ is up to $2\rho/\varepsilon \leq O(\sqrt{\log n}/\varepsilon)$ (but note that these edges are not actually stored). Thus, for each c, i, the total number of edges in $G_{c,i}$ crossing between different parts of partition $\mathcal{P}_{c,i}$ is at most $\tilde{O}(n/\varepsilon)$.

Once we have computed all the partitions $\mathcal{P}_{c,i}$, we make a second pass over the graph G, and finish storing the information needed for our estimation algorithm from Theorem 2.1. In particular, for every cut value $c \in \tilde{C}$ and class $i \in [l]$, we reproduce the stream of (sampled, reweighted) edges $G_{c,i}$ (using the same randomness as in the first pass). For each such stream $G_{c,i}$, we store data of two types. The first type is all edges that cross between the parts of $\mathcal{P}_{c,i}$. The second type is a sample of $1/\varepsilon$ edges going out from each vertex (among the edges that stay in that part), as required by the sketch from Theorem 2.1. Note that sampling $1/\varepsilon$ edges out from a vertex is easily accomplished using Reservoir Sampling [Vit85].

We remark that we need to use the same randomness to reproduce the exact same $G_{c,i}$ in both passes. Since Theorem 2.1 uses just second moment calculations, it is enough to use constant-wise independent random number generator for this purporse, which takes only an additional polylog(n) space.

After we are done with the second pass on G, our super-sketch actually contains $O(\log n)$ independent copies of the above sketch, each approximating every cut within factor $1+\varepsilon$ with high probability. This super-sketch estimates a cut value $w(T, \bar{T})$ by computing the $O(\log n)$ values obtained from $O(\log n)$ sketches, and reporting their median. By standard Chernoff bound, for each particular cut T, the new super-sketch fails to give a $1+\varepsilon$ approximation with probability at least $1 - 1/n^4$.

Finally, to evaluate the min-cut value, we enumerate over all $1 + 1/\log n$ approximate min-cuts of sparsifier H, and evaluate them using the super-sketch. We can use karger's algorithm [Kar00] to enumerate over all $1 + 1/\log n$ approximate min-cuts in $\tilde{O}(n^2)$ time, which is also the bound on the number of such approximate cuts. Thus there are only $\tilde{O}(n^2)$ queries to the super-sketch, and hence the super-sketch will success to output a $1+\varepsilon$ approximation on all of them, with probability at least $1 - \tilde{O}(n^2) \cdot 1/n^4 \geq 1 - 1/n$.

The total runtime is $\tilde{O}(n^4)$ since sketch evaluation for each set S takes $\tilde{O}(n^2)$ time. At the end of the first pass, we spend $\tilde{O}(n^3)$ time to compute the partitioning using the fast approximation algorithm for sparse-cut due to [She09]. □

D Sketch Size Lower Bound of $\Omega(n/\varepsilon)$

The following theorem proves that our sketch from Theorem 2.1 achieves optimal space up to polylogarithmic factors, even for unweighted graphs.

Theorem D.1. *Fix an integer n and $\varepsilon \in [2/n, 1/2]$. Suppose* **sk** *is a sketching algorithm that outputs at most $s = s(n, \varepsilon)$ bits, and* **est** *is an estimation algorithm, such that together, for every n-vertex graph G and subset $S \subset V$, with probability at least $9/10$ the estimation procedure is correct*

up to factor $1+\varepsilon$*, i.e.,*

$$\Pr\Big[\,\mathbf{est}(S,\mathbf{sk}(G)) \in (1\pm\varepsilon)\cdot|E(S,\bar{S})|\Big] \geq 9/10.$$

Then $s \geq \Omega(n/\varepsilon)$*.*

Proof. We will show how to encode a bit-string of length $l := n/(8\varepsilon)$ into a graph, so that, given its sketch $\mathbf{sk}(G)$, one can reconstruct any bit of the string with constant probability. Standard information-theoretical argument would then imply that $s \geq \Omega(l) = \Omega(n/\varepsilon)$.

Given a string $x \in \{0,1\}^l$, we embed it into a bipartite graph G on with $n/2$ vertices on each side, and vertex degrees bounded by $D := 1/(4\varepsilon)$ as follows. Partition the vertices on each side into disjoint blocks of D, and let the i-th block on the left side and on the right side form a (bipartite) graph which we call G_i, for $i = 1,\ldots,n/(2D)$. Then partition the string x in $n/(2D)$ blocks, each block is of length D^2 and describes the adjacency matrix of some bipartite G_i.

We now show that evaluating a bit from the string x corresponds to testing the existence of some edge (u,v) from some graph G_i, which we can do using the $1+\varepsilon$ approximating sketch only. Formally, let $\delta(S)$ be the cut value of the set S, i.e., $|E(S,\bar{S})|$, and observe that

$$\delta(\{u\}) + \delta(\{v\}) - \delta(\{u,v\}) = \begin{cases} 0 & \text{if } (u,v) \text{ is an edge in } G; \\ 2 & \text{otherwise.} \end{cases}$$

Since the considered values of $\delta(\cdot)$ are bounded by D, the sketch estimates each such value with additive error at most $\varepsilon D = 1/4$, which is enough to distinguish between the two cases. Furthermore, since we query the sketch only 3 times, the probability of correct reconstruction of the bit is at least $7/10$. The lower bound follows. □